

Specimen — illustrative example

Solstice Grid Services is a fictional organisation. This document comes from the scenario described on riskwoven.com/examples, and is not an assessment of any real entity. It is advisory output to support your own decisions: not legal advice, not a certification, and not an audit.

Written and reviewed by RiskWoven.com. Written by RiskWoven.com's AI drafting agent and checked by RiskWoven.com's independent review stage before release. The review is automated. It is not a human, legal or audit review.

AI Incident Response Procedure

Solstice Grid Services

Document type	procedure	Version	1.0 Draft
Owner	Head of Engineering	Customer decision role	Chief Technology Officer
Effective date	Upon customer release	Review cycle	Annual
Framework basis	RiskWoven AI practical best practices		
Jurisdiction	Ontario, Canada		
Prepared for	Solstice Grid Services	Industry	Energy and utilities
Organization size	101-500	Locations	Not specified

Advisory draft for accountable-owner and professional review. Validate legal, regulatory, contractual, technical and operational requirements before customer release. This is not legal advice, a certification, or a compliance determination.

1. Purpose

This procedure tells Solstice Grid Services staff exactly what to do when an AI system produces a harmful output, acts outside its defined boundary, exposes personal information, or

fails in a way that could affect grid operations or customers. It covers the autonomous grid-telemetry agents and the customer correspondence model.

This procedure sits beneath the AI Governance and Accountability Policy and the AI Risk Management Standard. Where an incident also meets the threshold for a security event, the Security Incident Response Procedure runs in parallel. Where personal information is exposed, the Privacy Breach Response Procedure runs in parallel. This procedure does not replace either of those instruments; it coordinates with them.

2. Scope

This procedure applies to:

- All AI systems operated by Solstice Grid Services, including the autonomous grid-telemetry agents and the outbound correspondence model.
- All staff, contractors and on-call engineers who operate, monitor or maintain those systems.
- All incidents, regardless of whether harm has been confirmed.

The procedure applies at all hours. The control room operates continuously. Engineering support runs Monday to Friday 07:00 to 19:00 Eastern, with on-call cover at all other times. Agents are halted rather than left unattended when no engineer is on call.

3. Definitions

AI incident. Any event in which an AI system at Solstice Grid Services produces an output, takes an action, or fails in a way that causes or could cause harm to grid stability, safety, customers, employees or the organisation. An incident includes boundary breaches, data exposure, model drift, harmful outputs and control failures, whether or not harm has been confirmed.

Harmful output. Any output from an AI system that causes or could cause physical harm, financial harm, reputational harm, or a breach of applicable law or Solstice Grid Services policy. For the grid-telemetry agents, this includes any action that affects grid stability or safety, any action outside the defined action boundary, and any action taken on stale or corrupted telemetry. For the correspondence model, this includes any message sent at scale that is inaccurate, misleading or discloses information it should not.

Model drift. A change in an AI system's behaviour over time, relative to its validated baseline, that has not been intentionally introduced. Drift may result from changes in input data distributions, infrastructure changes, or degradation of the model itself. Drift becomes an incident when it produces or could produce a harmful output.

Affected decision. Any operational action taken by an AI system, or any human decision made in reliance on an AI output, that may need to be reviewed, corrected or reversed because of an incident. For the grid-telemetry agents, this includes work orders opened, maintenance rescheduled and forecast inputs adjusted. For the correspondence model, this includes any outbound message sent to customers.

Containment. The set of actions taken to stop an AI system from causing further harm while the incident is investigated. For the grid-telemetry agents, containment means halting the agent. For the correspondence model, containment means suspending the send queue. Containment does not require harm to be confirmed; suspicion is sufficient.

4. Roles and Responsibilities

Incident Lead — Head of Engineering. Owns this procedure. Coordinates the incident response from detection through to closure. Makes containment decisions for Medium and High incidents. Escalates Critical incidents to the Chief Technology Officer immediately.

Executive Decision-Maker — Chief Technology Officer. Makes containment and remediation decisions for Critical incidents. Reviews every boundary breach within the target window set in Section 6, whether or not harm resulted. Holds review authority for this procedure.

On-Duty Engineer. The first technical responder. Receives reports, performs initial triage, executes containment steps, and preserves evidence. Reviews the agent action log daily and raises an incident for any action outside the defined boundary.

Control Room Supervisor. Assesses grid-operations impact. Confirms whether a halted agent has affected operational continuity. Coordinates with the on-duty engineer during Critical and High incidents.

Privacy Officer. Assesses whether an incident involves personal information. Determines whether the Privacy Breach Response Procedure must be activated. Manages any notification duties to affected individuals.

Chief Executive Officer. Leads all external communications: regulator notifications and customer communications. Acts on the advice of the Privacy Officer and the Chief Technology Officer.

5. Incident Severity Levels

Critical. Any of the following:

- An agent action that affects or could affect grid stability or safety.
- Any agent action taken outside the defined action boundary (switching load, touching protection settings, acting on the operational technology network).
- Exposure of customer consumption data linked to a service address.

High. Any of the following:

- An agent acting on stale or corrupted telemetry.
- A correspondence error sent at scale to customers.

Medium. A contained agent fault where no action was taken and no data was exposed.

Low. A monitoring gap with no operational effect and no data exposure.

Severity is assigned by the on-duty engineer at triage. The Incident Lead may revise severity at any point as new information emerges. Severity may only be downgraded with the agreement of the Incident Lead.

6. Response Targets

The following targets apply. They are set by Solstice Grid Services and are not derived from an external framework requirement.

Critical incidents:

- Acknowledgement and initial triage: within 15 minutes of report.
- Agent halted: immediately on suspicion, before triage is complete.
- Executive escalation to Chief Technology Officer: within 30 minutes of report; this is Solstice Grid Services' own organisational parameter and is not derived from an external requirement.
- Boundary breach reviewed by Chief Technology Officer: within 24 hours of detection, whether or not harm resulted.

High incidents:

- Acknowledgement and initial triage: on a target set by the Head of Engineering and recorded

in the AI Risk Management Standard before this procedure takes effect.

- Incident Lead notified: within the same window.

Medium incidents:

- Acknowledgement and initial triage: on a target set by the Head of Engineering and recorded in the AI Risk Management Standard before this procedure takes effect.

Low incidents:

- Logged and reviewed at the next scheduled engineering review.

The Head of Engineering must set and record the High and Medium acknowledgement targets in the AI Risk Management Standard before this procedure takes effect.

SPECIMEN

7. Step 1 — Detect and Report

Role: Any staff member, contractor or automated monitoring system.

7.1. Any person who observes or suspects an AI incident raises a report immediately. Reports may come from:

- The on-duty engineer's daily review of the agent action log.
- Automated monitoring alerts from the agent or correspondence system.
- A Control Room Supervisor observing unexpected grid behaviour.
- A customer or employee complaint.
- A third-party supplier notification.

7.2. The reporter contacts the on-duty engineer directly. Outside engineering support hours, the reporter uses the on-call contact. The on-duty engineer is the single point of receipt for all AI incident reports.

7.3. The reporter records the following at the time of report:

- Date and time of the observation.
- The AI system involved.
- A description of the observed behaviour or output.
- Any immediate operational effect they are aware of.

7.4. The on-duty engineer opens an incident record in the incident log immediately on receipt of the report. The incident record is the authoritative record for all subsequent steps.

8. Step 2 — Triage

Role: On-Duty Engineer, with Control Room Supervisor for grid-operations questions.

8.1. The on-duty engineer reviews the report and assigns an initial severity level using the criteria in Section 5.

8.2. The on-duty engineer records the following in the incident record:

- Initial severity level and the reason for that assignment.
- The AI system and the specific output, action or failure involved.
- The time of detection and the time of report.
- Whether the incident is ongoing.

8.3. If the incident is Critical or High, the on-duty engineer notifies the Incident Lead (Head of Engineering) immediately.

8.4. If the incident involves grid operations or the agent action log shows an action outside the defined boundary, the on-duty engineer also notifies the Control Room Supervisor immediately.

8.5. If the incident may involve personal information — including any exposure of customer consumption data — the on-duty engineer notifies the Privacy Officer immediately.

8.6. Triage does not delay containment. If the on-duty engineer suspects a Critical incident, containment in Step 3 begins before triage is complete.

9. Step 3 — Containment

Role: On-Duty Engineer. Incident Lead notified. Chief Technology Officer for Critical incidents.

9.1. **Grid-telemetry agents — Critical incidents.** The on-duty engineer halts the agent immediately on suspicion. Halting does not require confirmation of harm. The agent remains halted until a human clears it under Step 7.

9.2. **Grid-telemetry agents — High incidents.** The on-duty engineer restricts the agent to read-only mode or halts it, depending on the nature of the fault. The Incident Lead confirms the containment action within the High acknowledgement target.

9.3. **Correspondence model — any incident involving messages sent at scale.** The on-duty engineer suspends the send queue immediately. No further messages are sent until the queue is cleared by the Incident Lead.

9.4. The on-duty engineer records the containment action taken, the time it was taken, and who authorised it in the incident record.

9.5. The on-duty engineer preserves the agent action log, the inputs that produced the output or action, and any monitoring data from the period of the incident. Preservation means creating a read-only copy that cannot be altered. This evidence is retained in accordance with the Records Retention and Disposal Policy.

9.6. The on-duty engineer does not alter, delete or overwrite any log or input data related to the incident.

9.7. The Control Room Supervisor confirms whether the halted agent has affected operational continuity and advises the Incident Lead.

10. Step 4 — Escalation

Role: Incident Lead (Head of Engineering). Chief Technology Officer for Critical incidents.

10.1. For Critical incidents, the Incident Lead escalates to the Chief Technology Officer within the target window set in Section 6; that window is Solstice Grid Services' own organisational parameter and is not derived from an external requirement. The escalation includes:

- The incident record reference.
- The AI system involved.
- The containment action taken.
- The known or suspected impact on grid operations, customers or personal information.

10.2. The Chief Technology Officer decides whether additional containment is needed and whether external notification is required. The Chief Technology Officer records that decision in the incident record.

10.3. For any boundary breach — an agent action outside the defined action boundary — the Chief Technology Officer reviews the incident within 24 hours of detection, whether or not harm resulted. That review is recorded in the incident record.

10.4. Where the Privacy Officer has been notified, the Incident Lead coordinates with the Privacy Officer to determine whether the Privacy Breach Response Procedure must be activated. If it must, both procedures run in parallel from this point.

10.5. Where the incident may require regulator or customer notification, the Incident Lead briefs the Chief Executive Officer. The Chief Executive Officer leads all external communications. The timing and content of those communications are determined by the Chief Executive Officer on the advice of the Privacy Officer and the Chief Technology Officer, in accordance with applicable Ontario privacy law. The information supplied for this procedure does not specify which instruments apply; the Privacy Officer must confirm the applicable requirements.

10.6. Where a third-party AI supplier is involved, the Incident Lead initiates the Third-Party AI Review Procedure.

11. Step 5 — Impact Assessment

Role: Incident Lead, On-Duty Engineer, Privacy Officer, Control Room Supervisor.

11.1. The Incident Lead leads the impact assessment. The assessment identifies:

Affected data. The on-duty engineer identifies which data the AI system accessed or processed during the incident period. For the grid-telemetry agents, this includes telemetry inputs and any data used to generate the action. For the correspondence model, this includes the customer records used to generate messages. The Privacy Officer determines whether any of that data is personal information under applicable Ontario privacy law.

Affected outputs. The on-duty engineer lists every output or action the AI system produced during the incident period, using the preserved action log.

Affected decisions. The Incident Lead identifies every affected decision: work orders opened, maintenance rescheduled, forecast inputs adjusted, or correspondence sent. Each affected decision is recorded in the incident record.

Affected people. The Privacy Officer identifies the individuals whose personal information may have been involved. For customer consumption data, this may include accounts from the approximately 85,000 metered customers. The Privacy Officer records the estimated number of affected individuals.

Affected integrations and suppliers. The on-duty engineer identifies any downstream systems or third-party suppliers that received outputs from the AI system during the incident period.

11.2. The Incident Lead records the completed impact assessment in the incident record before moving to Step 6.

12. Step 6 — Correct Affected Decisions

Role: Incident Lead, On-Duty Engineer, Control Room Supervisor.

12.1. The Incident Lead reviews each affected decision identified in Step 5 and determines whether it can be corrected or reversed.

12.2. For each affected decision, the Incident Lead records in the incident record:
- Whether the decision is being corrected, reversed or left in place.

- The reason for that determination.
- The person responsible for executing the correction.
- The target date for completion.

12.3. The on-duty engineer executes corrections to agent actions using the reversal capability built into the agent action log. Every reversal is logged with the time, the engineer's identity and the incident record reference.

12.4. The Control Room Supervisor confirms that any correction to a grid-operations action has been applied correctly and that grid stability is not affected by the correction itself.

12.5. For correspondence sent at scale, the Incident Lead determines whether a correction or retraction must be sent to affected customers. The Chief Executive Officer leads any customer communication.

12.6. The Incident Lead records the outcome of each correction in the incident record, including the date completed.

13. Step 7 — Remediation and System Clearance

Role: On-Duty Engineer, Incident Lead, Chief Technology Officer.

13.1. The on-duty engineer identifies the root cause of the incident. Root cause analysis is recorded in the incident record.

13.2. The on-duty engineer proposes a remediation plan. The plan must address the root cause and must include:

- The specific change to be made to the AI system, its inputs, its boundary configuration or its monitoring.
- The person responsible for each change.
- The target date for each change.

13.3. The Incident Lead reviews the remediation plan and records that review in the incident record.

13.4. For Critical incidents, the Chief Technology Officer reviews the remediation plan before any halted agent is cleared to resume operation. That review is recorded in the incident record.

13.5. The on-duty engineer implements the remediation plan. Each completed change is recorded in the incident record with the date and the engineer's identity.

13.6. The on-duty engineer validates that the remediation has been applied correctly before the AI system resumes operation. Validation includes:

- Confirming that the root cause has been addressed.
- Confirming that the action boundary configuration is intact.
- Confirming that monitoring is active.
- Running the system in a supervised state for a period set by the Incident Lead and recorded in the incident record before returning to normal operation.

13.7. The Incident Lead records clearance to resume operation in the incident record. For Critical incidents, the Chief Technology Officer records that clearance.

14. Step 8 — Risk Assessment Update

Role: Incident Lead, Head of Engineering.

14.1. After remediation is complete, the Incident Lead updates the risk assessment for the affected AI system in the AI Risk Management Standard.

14.2. The update records:

- The incident reference.
- Whether the incident revealed a risk that was not previously identified or was rated lower than the incident warrants.
- Any change to the risk rating.
- Any new or revised control.

14.3. The Head of Engineering reviews the updated risk assessment and records that review.

15. Step 9 — Lessons Learned

Role: Incident Lead, all participants in the incident response.

15.1. The Incident Lead convenes a lessons-learned review after every Critical or High incident. The review takes place within a period set by the Head of Engineering and recorded in the AI Risk Management Standard before this procedure takes effect.

15.2. For Medium and Low incidents, the Incident Lead determines whether a lessons-learned review is warranted. That determination is recorded in the incident record.

15.3. The lessons-learned review considers:

- Whether detection was timely.
- Whether containment was effective.
- Whether the impact assessment was complete.
- Whether affected decisions were identified and corrected.
- Whether the remediation addressed the root cause.
- Whether this procedure needs to be updated.

15.4. The Incident Lead records the outcomes of the lessons-learned review in the incident record and in a separate lessons-learned log maintained by the Head of Engineering.

15.5. Where the review identifies a change to this procedure, the Head of Engineering initiates a procedure update under Section 17.

16. Incident Record Requirements

Every AI incident must have a single incident record that is opened at Step 1 and closed only when all steps are complete. The incident record must contain:

- Incident reference number.
- Date and time of detection and of report.
- AI system involved.
- Description of the output, action or failure.
- Severity level and the reason for that assignment.
- Containment action taken and the time it was taken.
- Impact assessment outcomes: affected data, outputs, decisions, people, integrations and suppliers.
- Correction and reversal outcomes for each affected decision.
- Root cause analysis.
- Remediation plan and completion dates.
- Validation outcomes.
- Risk assessment update reference.
- Lessons-learned outcomes.
- Date of closure and the name of the Incident Lead who closed it.

Incident records are retained in accordance with the Records Retention and Disposal Policy. The information supplied for this procedure does not specify a retention period; the Head of Engineering must confirm the applicable period with the Privacy Officer and record it in the AI Risk Management Standard.

17. Metrics

The Head of Engineering tracks the following metrics and reports them to the Chief Technology Officer on a cycle set by the Head of Engineering and recorded in the AI Risk Management Standard before this procedure takes effect.

- **AI incidents contained within target.** The number and percentage of incidents where containment was completed within the target window for the assigned severity level.
- **Affected decisions reviewed and corrected.** The number of affected decisions identified, the number reviewed, and the number corrected or reversed.

- **Percentage of required reviews completed on time.** The percentage of lessons-learned reviews, risk assessment updates and boundary-breach reviews completed within their target windows.
- **Exceptions past their expiry date.** The number of documented exceptions to this procedure that have passed their agreed expiry date without renewal or closure.
- **Material control failures and overdue corrective actions.** The number of incidents classified as Critical or High, and the number of corrective actions from those incidents that are past their target completion date.

18. Document Control

Document title: AI Incident Response Procedure

Document owner: Head of Engineering

Review authority: Chief Technology Officer

Review frequency: Annual, or following any Critical incident, whichever comes first.

The Head of Engineering reviews this procedure annually and after every Critical incident. The Chief Technology Officer holds review authority and records that review before any revised version takes effect. The review record is retained in accordance with the Records Retention and Disposal Policy.

This procedure is read alongside the AI Governance and Accountability Policy, the AI Risk Management Standard, the Security Incident Response Procedure, the Privacy Breach Response Procedure, and the Third-Party AI Review Procedure. Where those documents conflict with this procedure, the Head of Engineering raises the conflict with the Chief Technology Officer for resolution before the next incident response.

Owner review items

The drafting process identified the following in the supplied answers. Resolve each before customer release.

Type	Detail	Where
gap	The customer's answers specify response targets for Critical incidents (15-minute acknowledgement, 30-minute executive escalation, 24-hour boundary-breach review) but do not specify acknowledgement targets for High and Medium incidents. The Head of Engineering must set and record the actual values in the AI Risk Management Standard before this procedure takes effect.	Section 6 — Response Targets

gap	The customer's answers do not specify which Ontario privacy legislation or sector-specific regulation applies to Solstice Grid Services as an energy utility. The Privacy Officer must confirm the applicable instruments — candidates may include Ontario's privacy legislation and any sector-specific requirements from the Ontario Energy Board — and ensure that notification duties, timelines and individual rights obligations are reflected in the Privacy Breach Response Procedure and in the external communications process described in Step 4. No statutory contents have been reproduced in this procedu	Section 10 — Step 4, Escalation
gap	No retention period for incident records has been specified. This procedure directs the Head of Engineering to confirm the applicable period with the Privacy Officer and record it in the AI Risk Management Standard. The Records Retention and Disposal Policy governs retention; this procedure does not set a period independently.	Section 16 — Incident Record Requirements
gap	The lessons-learned review window for Critical and High incidents has been left as an organisational parameter. The Head of Engineering must set and record the actual value in the AI Risk Management Standard before this procedure takes effect.	Section 15 — Step 9, Lessons Learned
gap	The metrics reporting cycle has been left as an organisational parameter. The Head of Engineering must set and record the actual cycle in the AI Risk Management Standard before this procedure takes effect.	Section 17 — Metrics
gap	The customer's answers do not specify the physical or logical locations of the AI systems. If the grid-telemetry agents or the correspondence model are hosted by a third-party supplier, the Third-Party AI Review Procedure and the Third-Party Risk Management Policy will govern supplier notification and involvement. This procedure references the Third-Party AI Review Procedure at Step 4 but cannot specify supplier-specific contact details or contractual notification obligations because that information was not supplied.	Section 10 — Step 4, Escalation
vague	The customer's answers state that correspondence is reviewed before sending but do not specify who performs that review or what the review record looks like. This procedure treats the correspondence model as subject to the same incident response steps as the agents. If the pre-send review process has its own documented controls, the Head of Engineering should confirm that those controls are consistent with this procedure and that the review record is preserved as evidence.	Section 9 — Step 3, Containment
gap	No framework with verifiable control identifiers was supplied. This procedure is grounded in RiskWoven AI practical best practices and the customer's answers. No external control identifiers have been cited.	Entire document

Information the organization must supply

These values were not available from the intake and are required before this document can operate:

- Acknowledgement and containment targets for High incidents — to be set by the Head of Engineering and recorded in the AI Risk Management Standard before this procedure takes effect.

- Acknowledgement target for Medium incidents — to be set by the Head of Engineering and recorded in the AI Risk Management Standard before this procedure takes effect.
- Lessons-learned review window for Critical and High incidents — to be set by the Head of Engineering and recorded in the AI Risk Management Standard before this procedure takes effect.
- Metrics reporting cycle — to be set by the Head of Engineering and recorded in the AI Risk Management Standard before this procedure takes effect.
- Confirmation from the Privacy Officer of the applicable Ontario privacy legislation and any sector-specific regulatory requirements (candidates include Ontario's privacy legislation and Ontario Energy Board requirements), including notification timelines and individual rights obligations.
- Retention period for AI incident records — to be confirmed by the Head of Engineering with the Privacy Officer and recorded in the AI Risk Management Standard, consistent with the Records Retention and Disposal Policy.
- Confirmation of whether the grid-telemetry agents or the correspondence model are hosted by a third-party supplier, and if so, the supplier's contractual notification obligations and contact details for incident escalation.
- Confirmation that the pre-send review process for the correspondence model is documented and that its review records are preserved in a way that supports this procedure's evidence requirements.
- Reconciliation of this procedure with the Security Incident Response Procedure and the Privacy Breach Response Procedure to confirm that the parallel-running model described in Sections 1 and 10 is consistent with how those procedures define their own activation triggers.

Written and reviewed by RiskWoven.com. Written by RiskWoven.com's AI drafting agent and checked by RiskWoven.com's independent review stage before release. The review is automated. It is not a human, legal or audit review.